

WO 2004/054176

Shared medium communication system

The invention relates to a shared medium communication system comprising:

- a primary station arranged to receive an access request, process the access request and send a grant in response to the access request;
- a secondary station for sending the access request and for receiving the grant; and
- 5 - a shared medium coupling the primary station with the secondary station.

The invention also relates to a secondary station for sending an access request to a primary station and for receiving a grant from the primary station in response to the access request, the access request comprising a request for access to a shared medium.

10 The invention also relates to a primary station for receiving an access request from a secondary station, for processing the access request and for sending a grant to the secondary station in response to the access request, the access request comprising a request for access to a shared medium.

The invention also relates to a method of operating a shared medium communication system, the method comprising:

- 15
- a primary station receiving an access request, processing the access request and sending a grant in response to the access request,
 - a secondary station sending the access request to and receiving the grant from the primary station, the access request comprising a request for access to a shared medium.

20 The invention also relates to a method of sending an access request to a primary station and receiving a grant from the primary station in response to the access request, the access request comprising a request for access to a shared medium.

The invention also relates to a method of receiving an access request from a secondary station, processing the access request and sending a grant to the secondary station in response to the access request, the access request comprising a request for access to a 25 shared medium.

The invention also relates to a signal for use in a shared medium communication system.

An embodiment of the shared medium communication system described in the opening paragraph is known from the DOCSIS standard: "Data-Over-Cable Service Interface Specifications", Radio Frequency Interface Specification, SP-RFIv2.0-I02-020617 as issued by Cable Television Laboratories, Inc.

- 5 Each secondary station can send data traffic upstream to the primary station. The secondary station is a cable modem (CM) in a DOCSIS system, but may also be a set-top box, a television or a receiver. The primary station is a cable modem termination system (CMTS) in a DOCSIS system, but it may also be a head-end or a base station.

- 10 The data sent by all secondary stations is carried by the shared medium. When the data sent by two or more secondary stations collides on the medium, however, the data becomes unintelligible at the primary station.

- 15 A measure to prevent such a collision, which is described in the DOCSIS standard, is a request-grant mechanism. A secondary station that wants to use the medium sends an access request to the primary station. The primary station arbitrates between contending secondary stations and sends a grant downstream which indicates the secondary station that may use the medium. A secondary station may only send its data upstream after having received the grant. The request-grant mechanism prevents collisions between application data, but it does not prevent collisions between requests, which may be resolved with a collision resolution protocol. The request-grant mechanism is commonly applied in 20 shared medium communication systems including systems for media like coaxial or twisted pair cables, fiber, hybrid fiber coax (HFC) systems, and systems based on a passive optical network (PON). To make it fit for transmission over the medium, the access request is usually wrapped in a frame by functions in a lower layer of the ISO stack. This adds an overhead to each separate access request.

- 25 Each secondary station can handle several simultaneous service flows. A service flow can be seen as a unidirectional transport service of the secondary station. Handling a service flow by the secondary station can involve shaping, policing, and prioritizing the traffic in accordance with so-called quality of service traffic parameters defined for the service flow.

- 30 Each service flow generates its own access requests at the secondary station. Therefore, the access requests of all service flows contend for access to the medium. This deteriorates the performance of the communication system; especially it increases access delay.

The known shared medium communication system has a relatively low efficiency in handling the access requests and a relatively high access delay.

5 It is an object of the present invention to provide a shared medium communication system of the kind described in the opening paragraph, which has a relatively high efficiency in handling the access requests.

The object is realized in that the secondary station is arranged to merge several access requests into a multi-request and send the multi-request to the primary station, and in
10 that the primary station is arranged to receive the multi-request, process the multi-request and send the grants in response to the access requests merged in the multi-request.

15 The efficiency of the system in handling the access requests is improved because of the savings on the overhead of multiple separate access requests. Another advantage is that the access delay is reduced, because the merged access requests cannot mutually contend for access to the shared medium.

Advantageously, the access requests merged may also comprise a multi-request.

20 Advantageously, the multi-request is represented in a DOCSIS request frame by encoding it in the bits that indicate the so-called service ID (SID) and the request size,
being the number of requested mini-slots. All these bits, apart from a prefix, are available for
encoding the multi-request. At least one of the bits encodes the SIDs of the merged access
requests. At least one of the remaining bits encodes the request sizes of each of the merged
access requests. This has the advantages that ordinary DOCSIS request frames can be used.

25 Advantageously, the multi-request is made fit for transmission by putting it in a DOCSIS vendor-specific frame. This can be indicated in the FC_PARM field of a DOCSIS MAC frame by an arbitrary but specific bit combination, which can for example be included in a future standard. This has the advantage that there are more bits available for encoding the multi-request. This increases the number of access requests that can be efficiently merged in the multi-request. The so-called Req./Data transmit opportunity can be used for transmission
30 of vendor-specific MAC frames.

In an embodiment of the shared medium communication system according to the invention, the secondary station is arranged to adapt the merging of the access requests in dependence on histories of access previously merged, multi-requests previously sent and/or grants previously received.

In an embodiment of the shared medium communication system according to the invention, the secondary station is arranged to adapt the sending of the multi-request in dependence on histories of access requests previously merged, multi-requests previously sent and/or grants previously received.

5 Advantageously, the two above adaptations comprise combining parameters of at least two access requests that pertain to a single service flow. One example is summing the data sizes of at least two consecutive access requests for data transmittal. Another example for a slotted medium is summing the requested number of slots of at least two consecutive access requests. Yet another example inspired on the DOCSIS standard, is summing the
10 requested number of so-called mini-slots of at least two request frames from a single service flow.

Advantageously, these adaptations exploit the trade-off between the potential gains in merging two particular access requests and the disadvantages of an actual collision. If, on the one hand, there is a relatively high traffic load on the system, it pays off to merge a
15 relatively high number of access requests into the multi-request. If, on the other hand, there is a relatively low traffic load on the system, there is relatively little gained because access requests are seldomly merged. The actual traffic load can be estimated by deriving from the histories.

20 Advantageously, the primary station obtains the access requests contained in the multi-request and processes each access request as if it was received separately.

Advantageously, the primary station uses a means for processing access requests partially or as a whole, also for processing the multi-request.

25 The above object and features of the present invention will be more apparent from the following description of the preferred embodiments with reference to the drawings wherein:

Fig. 1 shows a block diagram of an embodiment of a communication system according to the invention.

30 Fig. 2 shows a multi-request for use in an embodiment of a communication system according to the invention.

Fig. 3 shows some graphs illustrating the performance of a communication system according to the invention.

Fig. 1 shows a block diagram of an embodiment of a communication system according to the invention. A shared medium communication system 100 comprises a primary station 102 and at least one secondary station 104. The primary station may be a cable modem termination system (CMTS), but it may also be a head-end, a concentrator, a management center or a base station. The secondary station may be a cable modem (CM), but it may also be a set-top box, a television, a network terminator, a receiver or so-called customer premises equipment.

A shared medium 106 couples the primary station 102 with each secondary station 104. The medium and the network usually have a tree topology, but also bus or star topologies are possible. The medium may be coaxial or twisted pair cables, fibers, hybrid fiber coax (HFC) or a passive optical network (PON).

To avoid collisions on the shared medium 106, each secondary station 104 has to send a request 108 to the primary station 102. Only after receiving a grant 110, the secondary station 104 can send further data to the primary station 102. The further data may comprise any content, like for example audio and/or video streams, data files, transactions, subscriber feedback, other broadcasts. Instead of sending several separate requests 108, a secondary station 104 can merge the requests 108 into a multi-request 108 and send the multi-request 108 to the primary station 102. After receiving the multi-request 108, the primary station 102 processes the multi-request 108 and sends grants 110 in response to the merged requests in the multi-request 108.

Fig. 2 shows a multi-request for use in an embodiment of a communication system according to the invention. The multi-request is represented in a DOCSIS request frame (RF) 200. A request frame is typically composed of a payload and a header.

The payload of a DOCSIS request frame 200 consists of a 14-bit SID 210 and an 8-bit request size 212. The header (not shown) includes physical-layer overhead such as:

- a guard band for safely separating subsequent bursts from different secondary stations;
- a preamble that is required for demodulation of the received burst;

and some MAC-layer overhead, such as:

- a control field for identifying a request as such and
- a check sequence for error detection.

As a result, only a relatively small part of the request frame is dedicated to the payload.

According to the invention, a special SID management scheme and a request-partitioning rule can be used as detailed in the following description.

The SID 210 is further distinguished in a prefix (PF) 220 of L bits, a secondary station number (SSN) 222 of N bits and requesting service flows (SFs) 224 which 5 is a field of the remaining M bits. For a DOCSIS request frame 200, $M=14-L-N$.

Predefined prefixes 220, for example '01' and '001', are used to indicate brand secondary stations or other secondary stations capable of merging access requests. The primary station treats a brand secondary station in a special way. So-called special SIDs (SSIDs) start with one of these prefixes 220. All other SIDs are called normal SIDs (NSIDs).

10 The M bits of the requesting service flows 224 can be used to create a set of M SSIDs for the secondary station indicated by the SSN 222. Each M bit represents a service flow. A bit is set to indicate a merged request for the respective service flow. If multiple bits are set, the multi-request contains requests for each of the respective service flows. The maximum number of set bits is bounded from above by R, which is in turn bounded by 8 for 15 a DOCSIS frame 200. For each bit pattern of the requesting service flows 224, a structuring of the request size 212 into a number of request subfields is predefined. Consequently, with r bits set in the requesting service flows 224, the 8 bits are subdivided into r successive subfields of size $f_{r1}, f_{r2}, \dots, f_{rr}$ respectively, with for any r the sum over i of f_{ri} being 8. Subfield i encodes the request size of the merged request for the service flow of the ith set bit 20 in the requesting service flows 224.

In a concrete example, the prefix used to indicate a brand secondary station is the bit string '01', with $L=2$. The secondary station number is encoded in the next 8 bits, so $N=8$, implying a maximum of 256 brand secondary stations for this prefix. The remaining $M=14-2-8=4$ bits are available for encoding the service flows 224. Taking $R=2$, so that at 25 most 2 requests are merged in the multi-request, the subdivision of the original request field can be, for example : $f_{11} = 8$, and $f_{21} = f_{22} = 4$. That is, if only one SSID is represented in the combined SSID, then its request size can be up to 255 (8 bits), whereas, if there are 2 SSIDs represented in the combined SSID, each of them has a maximum request size of 15 (4 bits). As an example of the latter case, if the multi-request contains a SID 210 equal to the bit 30 string 01 11001011 0101 and a request size 212 equal to the bit string 1001 0111, then the request for SSID 01 11001011 0100 has request size 1001 (= 9) and the request for SSID 01 110010110 0001 has request size 0111 (= 7).

People skilled in the art can think of alternative ways to encode multi-requests in a request frame.

The number of available SSIDs has an upper limit for each secondary station. Therefore, the primary station executes some logic to determine if a SSID or a NSID is assigned upon the creation of a new service flow at a secondary station.

A secondary station may have separate sets of SSIDs assigned to it, provided
5 that it treats them independently, just like the simultaneous presence of more than one NSID. One possibility to implement this at the brand secondary station is to give SSIDs priority over NSIDs, each of them being handled in FCFS order. Alternatively, the access request and the multi-requests could be served in FCFS, a round robin, or otherwise fair manner. Finally, it is possible that SSIDs are also assigned to non-brand secondary stations, but the primary station
10 should consider these as NSIDs.

Merging requests generally reduces the maximum request size of the merged requests. One option to alleviate the impact is to use piggybacking. A first request only requests a small number of mini-slots to transmit only a part of the frame. The request for the next part of the frame is attached to the previous part of the frame.

15 Figure 3 shows some graphs illustrating the performance of a communication system according to the invention. In this case the system consists of 100 secondary stations and has 5 service flows at each secondary station. Shown is the mean of the connection active time (CAT) as a function of the traffic load on the system for several scenarios.

20 The CAT is defined here as the time difference (in units of time) between the instant that a first request is generated at a secondary station, until the instant that this request (possibly as part of a multi-request if merged with other requests) is successfully received at the primary station. In Figure 3, one unit of time corresponds to the duration of a single-request frame and an access request takes up one third of such frame.

25 The curves that are drawn in Figure 3 show the mean CAT without merging (sc. 1) and with merging (sc. 2), the latter for $R = 1, 2, 3, 4$ and 5 , as indicated in the legend of Figure 3. In addition, two curves depict analytical results (anal.) that validate the simulation results. For a relatively high traffic load, the mean CAT decreases significantly when merging is applied and when R increases. Only merging with $R = 1$ approaches the relatively large mean CAT of not merging (sc. 1). This is no surprise, as merging with $R = 1$
30 corresponds to servicing the service flows of an individual secondary station in FCFS order.

For a relatively low traffic load however, the reverse effect is observed: the mean CAT increases significantly when R increases.

This implies that making R dependent on the traffic load pays off in terms of a lower mean CAT, with R = 1 for a relatively low traffic load and increasing R to 4 or 5 towards a relatively high traffic load.